

Listing of the Claims

This listing of claims will replace all prior versions and listings of claims in the application.

Listing of Claims:

1. (currently amended) A method of unsupervised training of acoustic models of a segmentation-based automatic speech recognition system, comprising ~~the steps of:~~

receiving correct segment-based alignment data that represents a correct alignment of ~~an~~ a first sequence of utterance features ~~that was received by the speech recognition system;~~

receiving ~~wrong~~ incorrect segment-based alignment data that represents an incorrect alignment of ~~the~~ a second sequence of utterance features ~~that is known to be incorrect based on information received from by the speech recognition system and describing the utterance;~~

identifying a first phoneme in the correct alignment data that corresponds to a second phoneme in the ~~wrong~~ incorrect alignment data ~~and in which the first phoneme received a worse recognizer score than the second phoneme; and~~

modifying a first acoustic model of the first phoneme by moving at least one mean value thereof closer to ~~the~~ corresponding feature values ~~used to score the first phoneme in the first sequence of utterance features.~~

2. (currently amended) A method as recited in Claim 1, further comprising ~~the steps of:~~

modifying a second acoustic model of the second phoneme by moving at least one mean value thereof farther from ~~the~~ corresponding feature values ~~used to score the second phoneme in the second sequence of utterance features.~~

3. (currently amended) A method as recited in Claim 1, wherein ~~receiving the~~ correct alignment data ~~comprises the step of receiving correct alignment~~ includes data that represents a segment alignment of a less than highest scoring hypothesis from among n-best hypotheses of an utterance ~~that was received by the speech recognition system.~~

4. (currently amended) A method as recited in Claim 1, wherein ~~receiving wrong~~ the incorrect alignment data ~~comprises the steps of receiving wrong alignment data that represents an alignment of the utterance that is known to be incorrect based on user confirmation information received from the speech recognition system in response to prompting a speaker to confirm the utterance.~~

5. (currently amended) A method as recited in Claim 1, wherein ~~receiving the~~ correct alignment data ~~comprises the steps of receiving correct alignment data that represents an alignment of the utterance that is known to be correct based on user confirmation information received from the speech recognition system in response to prompting a speaker to confirm the utterance.~~

6. (currently amended) A method as recited in Claim 1, further comprising: ~~the step of iteratively repeating the identifying and modifying steps for all phonemes in the correct alignment data that correspond to one or more phonemes in the wrong alignment data.~~

7. (currently amended) A method as recited in Claim 2, further comprising: ~~the step of iteratively repeating the identifying and modifying steps for all phonemes in the wrong alignment data that correspond to one or more phonemes in the correct alignment data.~~

8. (currently amended) A method as recited in Claim 1, wherein ~~the step of moving at least one mean value farther from the~~ closer to corresponding feature value used to score it comprises values in the first sequence of utterance features includes subtracting a multiple of the corresponding feature value values from the at least one mean value of the first acoustic model.

9. (currently amended) A method as recited in Claim 1, wherein ~~the step of moving at least one mean value farther from the~~ closer to corresponding feature value used to score it comprises values in the first sequence of utterance features includes modifying the at least one mean value of the first acoustic model by approximately 2%.

10. (currently amended) A method as recited in Claim 1, wherein ~~the first acoustic model includes a plurality of model components and wherein~~ modifying a first acoustic model further ~~comprises the steps of~~ includes modifying a set of model components associated with the first phoneme by moving ~~all the~~ each mean ~~values~~ value thereof closer to the corresponding feature ~~values used to score the phoneme in the first sequence of utterance features.~~

11. (currently amended) A method as recited in Claim 2, wherein ~~the second acoustic model includes a plurality of model components and wherein~~ modifying the second acoustic model further ~~comprises the steps of~~ includes modifying a set of model components associated with the second phoneme by moving ~~all the~~ each mean ~~values~~ value thereof ~~closer to~~ farther from the corresponding feature values ~~used to score the phoneme in the second sequence of utterance features.~~

12. (currently amended) A method of improving performance of a segmentation-based automatic speech recognition system (ASR) ~~by training its acoustic models using information obtained from a particular application in which the ASR is used, comprising the steps of:~~

receiving a correct segment-based alignment of ~~an~~ a first sequence of utterance features ~~that was received by the ASR;~~

receiving an incorrect segment-based alignment of ~~the~~ a second sequence of utterance features ~~that is known to be incorrect based on information received from by the speech recognition system~~ ASR in the context of ~~the~~ a particular application using the ASR;

identifying a first phoneme in the ~~known~~ correct segment-based alignment that corresponds to a second phoneme in the incorrect segment-based alignment;

modifying a first acoustic model of the first phoneme by moving at least one mean value thereof closer to corresponding feature values ~~used to score the first phoneme in the first sequence of utterance features.~~

13. (currently amended) A method as recited in Claim 12, further comprising ~~the steps of:~~

modifying a second acoustic model of the second phoneme by moving at least one mean value thereof farther from ~~the~~ corresponding feature values ~~used to score the second phoneme in~~ the second sequence of utterance features.

14. (currently amended) A computer-readable medium carrying one or more sequences of instructions for training acoustic models of a segmentation-based automatic speech recognition system, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:

receiving correct segment-based alignment data that represents a correct segment alignment of ~~an~~ a first sequence of utterance features ~~that was received by the speech recognition system;~~

receiving ~~wrong~~ incorrect segment-based alignment data that represents an incorrect alignment of ~~the~~ a second sequence of utterance features ~~that is known to be incorrect based on information received from by the speech recognition system and describing the utterance;~~

identifying a first phoneme in the correct alignment data that corresponds to a second phoneme in the ~~wrong~~ incorrect alignment data; and

modifying a first acoustic model of the first phoneme by moving at least one mean value thereof closer to corresponding feature values ~~used to score the first phoneme in the first~~ sequence of utterance features.

15. (currently amended) A computer-readable medium as recited in Claim 14, wherein the instructions further comprise instructions for carrying out the steps of:

modifying a second acoustic model of the second phoneme by moving at least one mean value thereof farther from ~~the~~ corresponding feature values ~~used to score the second phoneme in~~ the second sequence of utterance features.

16. (currently amended) A segmentation-based automatic speech recognition system ~~that provides improved performance by training its acoustic models according to information about an application with which the system is used,~~ comprising:

a speech recognizer that includes one or more processors;
non-volatile storage coupled to the speech recognizer and comprising a plurality of segmentation alignment data and a plurality of acoustic models;
a computer-readable medium coupled to the speech recognizer and carrying one or more sequences of instructions for training acoustic models, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to perform the steps of:
receiving correct segment-based alignment data that represents a correct segment alignment of ~~an~~ a first sequence of utterance features ~~that was received by the speech recognition system;~~
receiving ~~wrong~~ incorrect segment-based alignment data that represents an incorrect alignment of ~~the~~ a second sequence of utterance features ~~that is known to be incorrect based on information received from~~ by the speech recognition system ~~and describing the utterance;~~
identifying a first phoneme in the correct alignment data that corresponds to a second phoneme in the ~~wrong~~ incorrect alignment data; and
modifying a first acoustic model of the first phoneme by moving at least one mean value thereof closer to corresponding feature values ~~used to score the first phoneme in the first sequence of utterance features.~~

17. (currently amended) A speech recognition system as recited in Claim 16, wherein the instructions further comprise instructions for carrying out the steps of:

modifying a second acoustic model of the second phoneme by moving at least one mean value thereof farther from ~~the~~ corresponding feature values ~~used to score the second phoneme in the second sequence of utterance features.~~

18.-20. canceled.